TRIPS: A Distributed Explicit Data Graph Execution (EDGE) Microprocessor

Madhu Saravana Sibi Govindan,
Doug Burger, Steve Keckler,
and the TRIPS Team

Hot Chips 19, August 2007
Computer Architecture and Technology Laboratory
Department of Computer Sciences
The University of Texas at Austin
www.cs.utexas.edu/users/trips

Recent Trends

- Scaling challenges for conventional superscalar processors
  - Power and pipeline limits impede clock rate growth
  - Wire delays cause overheads for concurrency
  - Complexity of large monolithic architectures
- Industry shift to multicore architectures
  - Work well for certain types of workloads
  - Big challenges: how to program/Amdahl's law
  - Single-thread performance is still important

TRIPS – A Technology Scalable Architecture

- Goals
  - High single-thread performance through ILP
  - Exploit concurrency at multiple granularities
  - Scalable with technology trends
- Key technologies
  - Explicit data graph execution (EDGE) ISA
  - Distributed processor microarchitecture
  - Distributed non-uniform (NUCA) L2 cache
  - Tiled and networked design
- Hot Chips 2005 - Chip in RTL design phase
- Hot Chips 2007 - Chip/system complete
  - Manufacturing/bring-up complete
  - Performance tuning in progress
  - This talk focuses on chip and system implementation

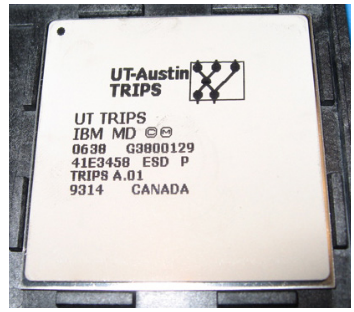

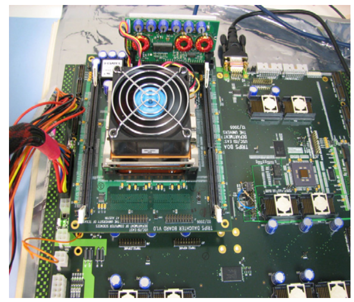

Outline

- Explicit Data Graph Execution (EDGE) ISAs
  - ISA support for distributed execution
- TRIPS networked microarchitecture
  - Processor microarchitecture
  - Non-uniform cache memory system
- TRIPS chip and system implementation
  - Custom ASIC chip and system boards
  - Preliminary performance results

Explicit Data Graph Execution (EDGE)

- Two key features
  - Program broken into sequence of instruction blocks
    - Execution model: fetch, execute, and commit blocks atomically
    - Amortize overheads over many instructions
  - Within a block, instructions explicitly encode producer-consumer communication
    - Producer instructions explicitly targets consumers and send results directly to them (without using a register file)
    - Instructions “fire” when all operands arrive
    - Any instruction may be predicated (conditionally executed)
- TRIPS blocks - up to 128 instructions
  - Compile-time techniques to create large blocks
  - Average block sizes typically greater than 45 instructions
  - Long-term goal: hide hard-to-predict branches inside blocks

TRIPS Execution Model

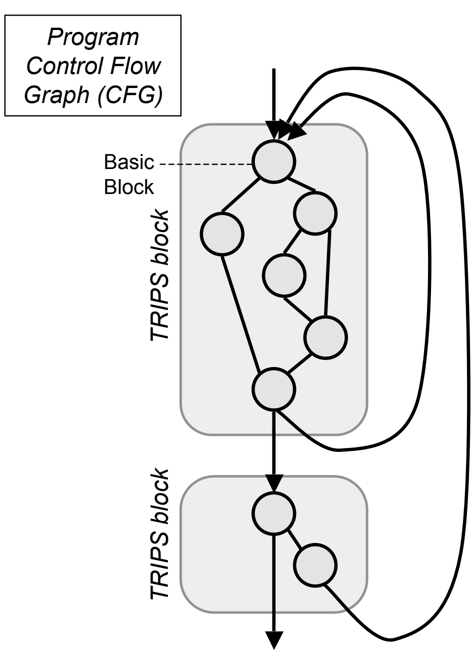

CFG Loop Example

if-then-else example

```
if (p==0)
  z = a * 2 + 3;
else
  z = b * 3 + 4;
```

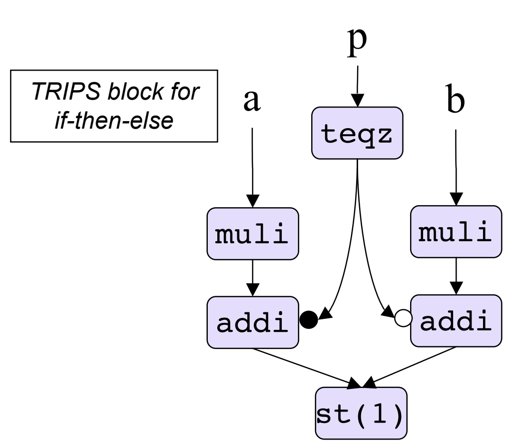

if-then-else Example

TRIPS Execution Model

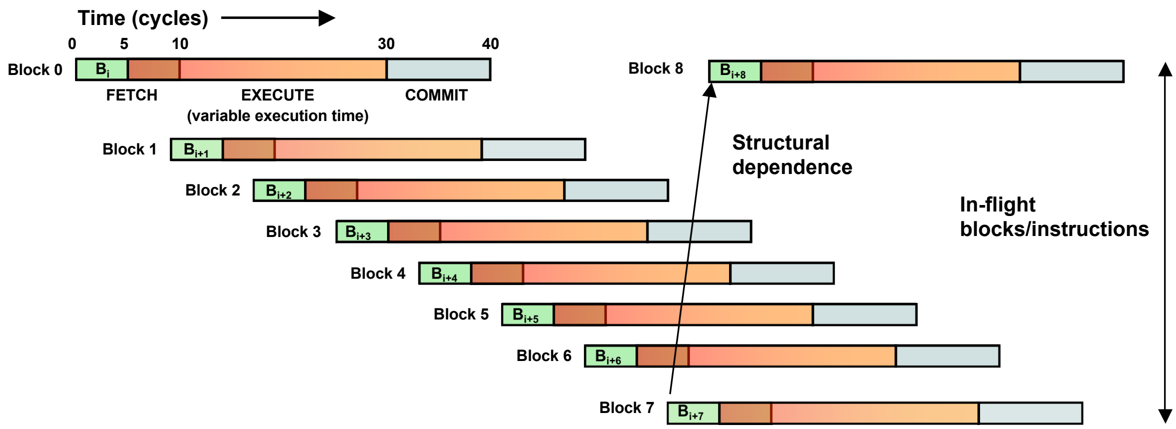

- Fetch/Execute/Commit overlapped across multiple blocks
  - Can execute up to 8 blocks at a time via speculation
- Exposes concurrency with a very large instruction window
  - Single-threaded mode with up to 1024 instructions
  - Simultaneous multi-threaded (SMT) mode with up to 4 threads and 256 instruction window per thread

TRIPS Prototype Chip

- 2 TRIPS Processors
- NUCA L2 Cache
  - 1 MB, 16 banks
- On-Chip Network (OCN)
  - 2D mesh network
  - Replaces on-chip bus
- Controllers
  - 2 DDR SDRAM controllers
  - 2 DMA controllers
  - External Bus Controller (EBC)
    - Interfaces with PowerPC 440GP (control processor)
  - Chip-to-Chip (C2C) network controller
- Clocking
  - 2 PLLs
  - 4 Clock domains
    - 1x and 2x SDRAM
    - Main and C2C
  - Clock tree
    - Main domain has 4 quadrants to limit local skew

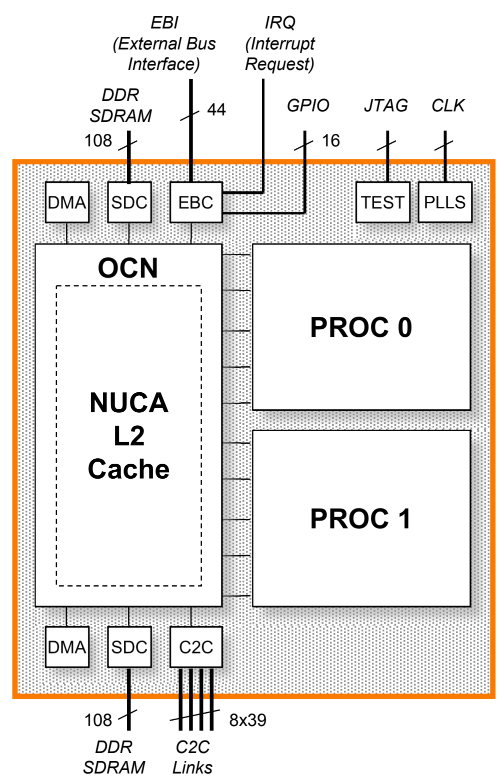

TRIPS Microarchitecture Principles

- Distributed and tiled architecture
  - Small and simple tiles (register file, data cache bank, etc.)
  - Short local wires
    - Tiles are small: 2-5 mm^2 per tile is typical
  - No centralized resources
- Networks connect the tiles
  - Networks implement distributed protocols (I-fetch, bypass, etc.)
    - Includes well-defined control and data networks
  - Networks connect only nearest neighbors
  - No global wires
- Design modularity and scalability
  - Design productivity by replicating tiles (design reuse)
  - Networks extensible, even late in design cycle

TRIPS Tile-level Microarchitecture

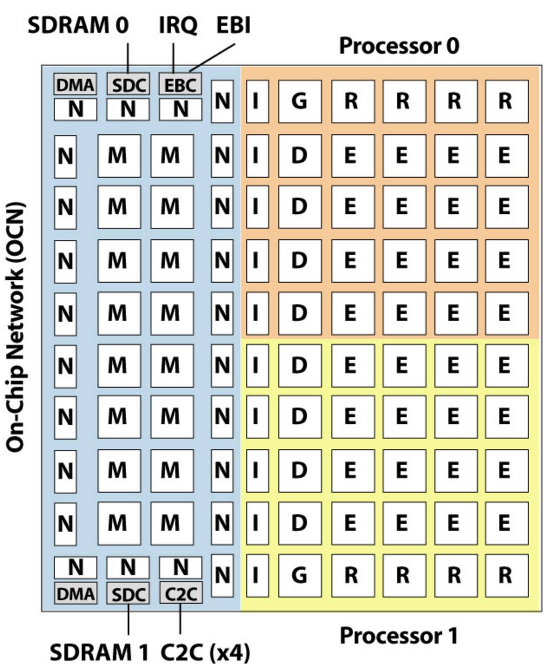

TRIPS Tiles

G: Processor control - TLB w/ variable size pages, dispatch, next block predict, commit

R: Register file - 32 registers x 4 threads, register forwarding

I: Instruction cache - 16KB storage per tile

D: Data cache - 8KB per tile, 256-entry load/store queue, TLB

E: Execution unit - Int/FP ALUs, 64 reservation stations

M: Memory - 64KB, configurable as L2 cache or scratchpad

N: OCN network interface - router, translation tables

DMA: Direct memory access controller

SDC: DDR SDRAM controller

EBC: External bus controller - interface to external PowerPC

C2C: Chip-to-chip network controller - 4 links to XY neighbors

TRIPS Processor

- EDGE ISA blocks mapped to array of tiles
  - Compile-time scheduler decides where (not when) instructions execute
- TRIPS: aggressive processor capabilities
  - Up to 16 instructions per cycle
  - Up to 4 loads/stores per cycle
  - Up to 64 outstanding L1 data cache misses
  - Up to 1024 dynamically scheduled instructions
  - Up to 4 simultaneous multithreading (SMT) threads
- Memory system
  - 4 simultaneous L1 cache fills per processor
  - Up to 16 simultaneous L2 cache accesses
  - Up to 16 outstanding L2 cache misses

Non-Uniform (NUCA) L2 Cache

- 1MB L2 cache
  - Sixteen tiled 64KB banks
- On-chip network
  - 4x10 2D mesh topology
  - 128-bit links, 366MHz (4.7GB/sec)
  - 4 virtual channels prevent deadlocks
  - Requests and replies are wormhole-routed across the network
- Up to 10 memory requests per cycle
- Up to 128 bytes per cycle returned to the processors
- Individual banks reconfigurable as scratchpad

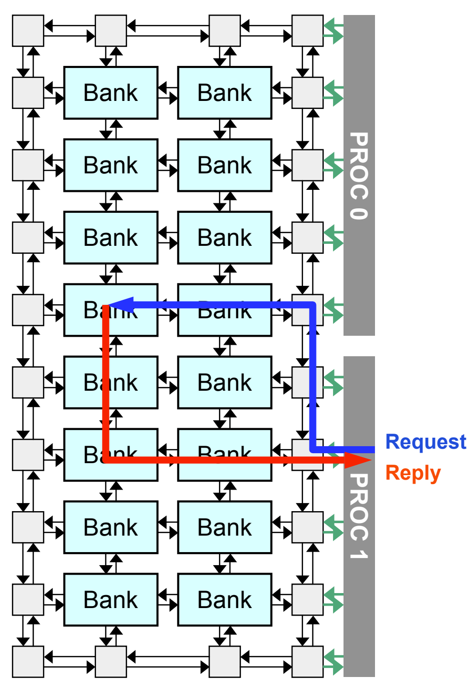

TRIPS Project Timeline

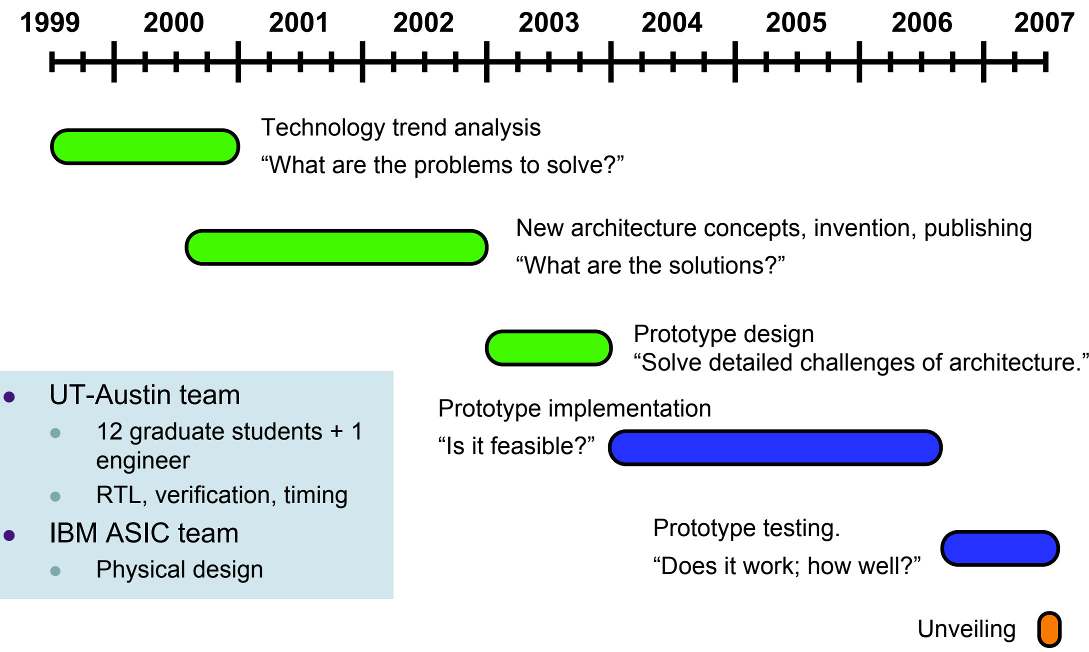

- UT-Austin team
  - 12 graduate students + 1 engineer
  - RTL, verification, timing
- IBM ASIC team
  - Physical design

TRIPS Chip Implementation

Process Technology	130nm ASIC with 7 metal layers
Die Size	18.3mm x 18.37mm (336 mm^2)
Package	47mm x 47mm BGA
Pin Count	626 signals, 352 Vdd, 348 GND
# of placed cells	6.1 million
Transistor count (est.)	170 million
# of routed nets	6.5 million
Total wire length	1.06 km
Power (measured)	36W at 366MHz (chip has no power mgt.)
Clock period	2.7ns (actual) 4.5ns (worse case sim)

Die Photos

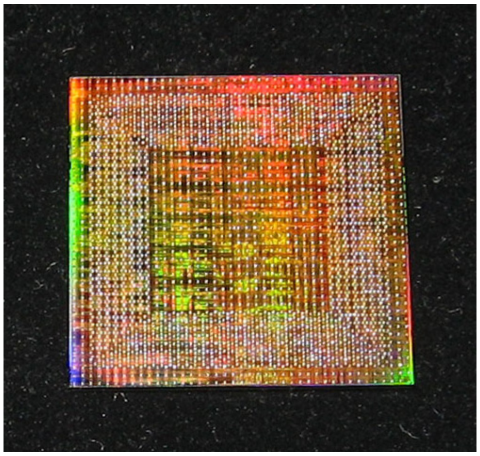

With C4 Array

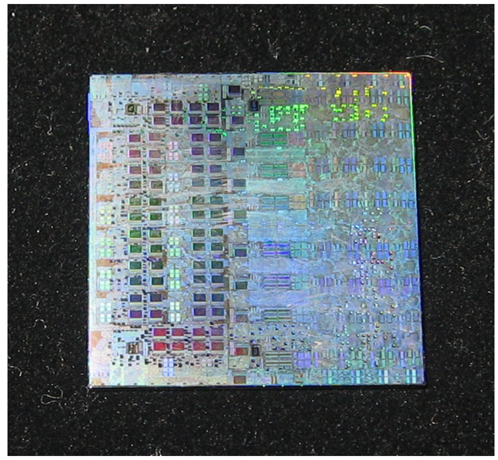

Without C4 Array

Benefits of Tiled Design

- Design modularity
  - 11 different tiles, instantiated a total of 106 times
  - Clean interfaces at tile boundaries
- Verification - no hardware bugs
  - Tiles verified extensively before stitching together
- Place and Route - hierarchical in nature
  - Wiring only between nearest neighbors
  - But - each physical instance was a little different
- Timing - trivial at top level, communication planned
  - No global wires or timing paths

TRIPS Daughtercard

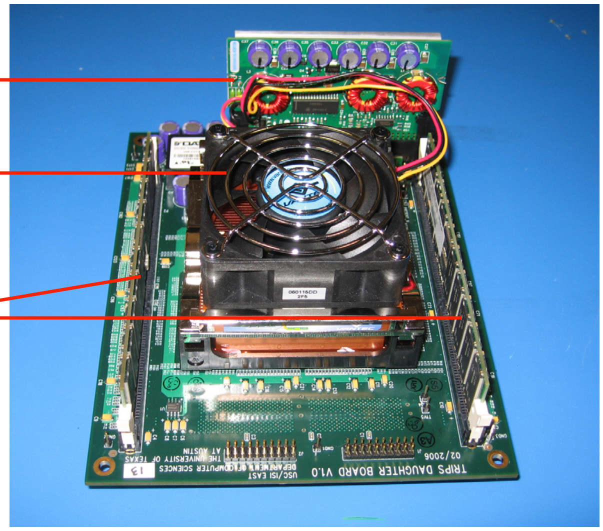

- 12V to daughtercard, stepped down to 1.5V, 2.5V, and 3.3V
- 12 GFlops peak
- 45 W total (worst case)

TRIPS Motherboard

- 1 motherboard includes:
  - 4 daughter-boards
  - 4 TRIPS chips
  - 8 GBytes DRAM
  - PowerPC 440GP control processor
  - I/O: ethernet, serial, C2C links
  - FPGA I/O interface
- Peak performance
  - 48 GFlops at 366 MHz
  - 180 Watts

TRIPS Multi-board System

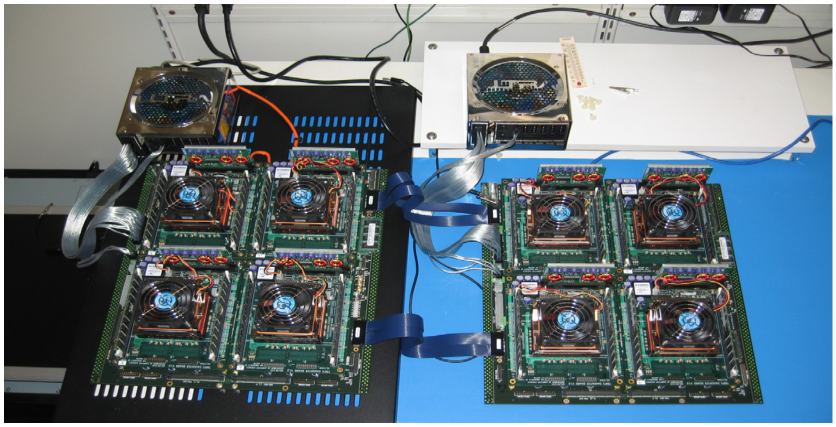

- Extensible to 8 boards (64 processors)
  - Micro-coax cables extend C2C network across boards
- Peak performance
  - 380 GFlops, 1.4KW
- Parallel message-passing software

TRIPS Prototype System Layers

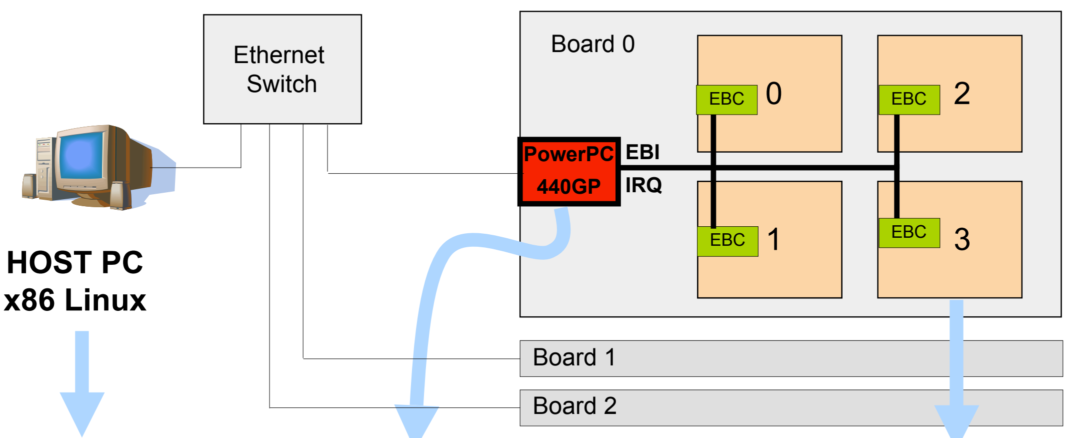

- *TRIPS Resource Manager* (TRM)
- File system
- Runtime services
- Login/debug/etc.

- Local Resoure Manager (LRM) listens to HostPC
- Runs embedded Linux
- PowerPC EBI device driver to control TRIPS chips
- PowerPC EBI↔TRIPS EBC

- Runs TRIPS apps
- Interrupts PowerPC if necessary
- System calls, exceptions

Preliminary Performance (HW)

- Challenges
  - Different technology and ISAs
  - Different processor-to-memory clock ratio
  - TRIPS compiler fine-tuning in progress
- Cycle-to-cycle comparison on multiple HW platforms
  - TRIPS Performance counters
  - PAPI - Performance API on Linux systems for others
- Applications
  - Compiled + hand-optimized
    - Mix of kernels and full algorithms
  - Compiled only
    - The Embedded Microprocessor Benchmark Consortium (EEMBC)
    - Versabench (MIT)
  - SPEC benchmarks in progress

Processor	Clock Speed	Memory Speed	Process Technology
TRIPS	366 MHz	200 MHz DDR	130 nm
Core 2	1.6 GHz (underclocked)	800 MHz DDR2	65 nm
Pentium 4	3.6 GHz	533 MHz DDR2	90 nm
Pentium 3	450 MHz	100 MHz SDRAM	250 nm

TRIPS vs. Conventional Processors: Kernels

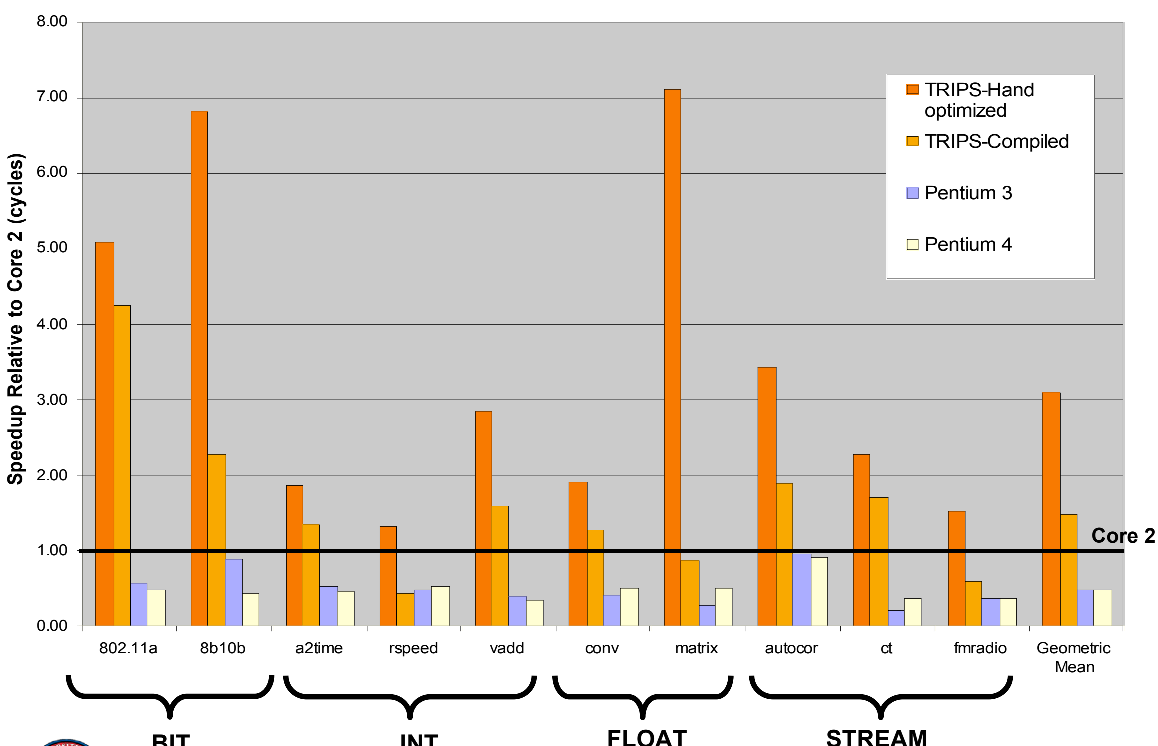

TRIPS vs. Conventional Processors
EEMBC and signal processing (compiled)

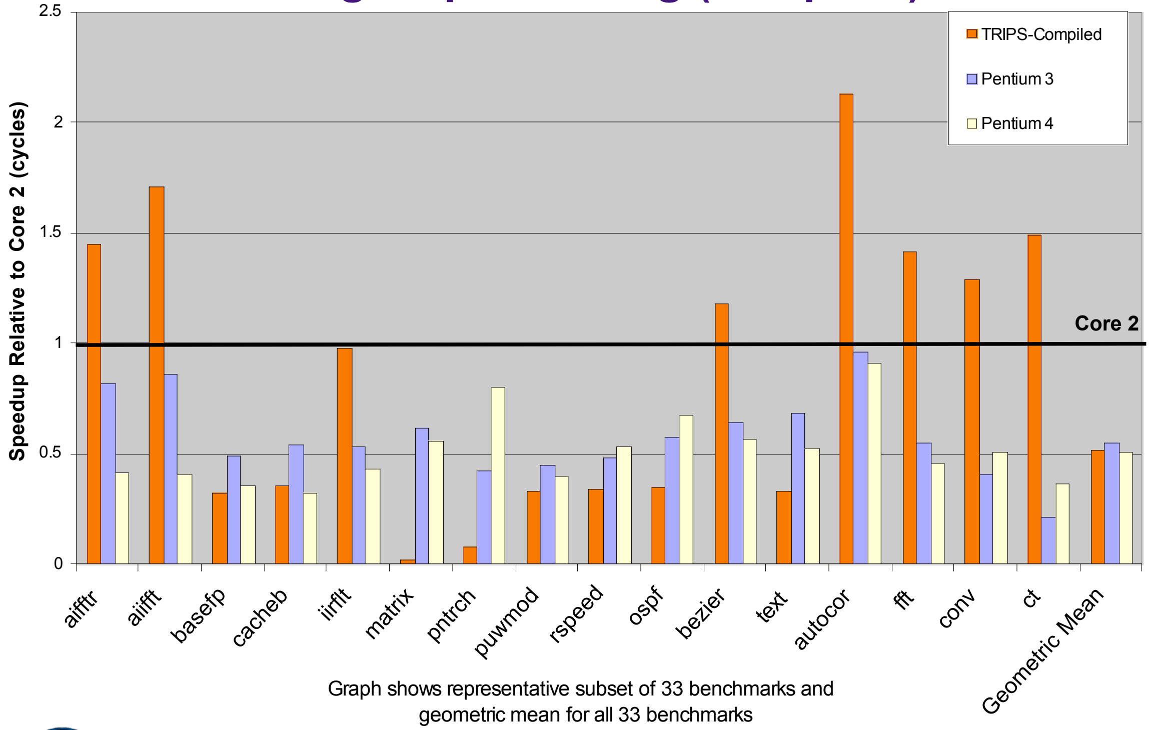

Graph shows representative subset of 33 benchmarks and geometric mean for all 33 benchmarks

Summary

- TRIPS prototype demonstrates feasibility of:
  - Explicit Data Graph Execution (EDGE) ISAs
  - Distributed processor and memory microarchitectures
  - Scaled and tiled uniprocessors
  - Non-uniform cache architectures (NUCA)
  - Recompilation required, but no change to source code
- Performance is promising
  - First generation TRIPS is 3.1x faster for hand-optimized code
  - Compiled-code reasonable and improving
- Identified several opportunities for improvements
  - Enhanced operand networks
  - Dynamic aggregation of processor cores
- TRIPS-like architectures are a potential “alternative” to multicore

Acknowledgements

- TRIPS Hardware Team
  - Raj Desikan, Saurabh Drolia, Madhu Sibi Govindan, Divya Gulati, Paul Gratz, Heather Hanson, Changkyu Kim, Haiming Liu, Robert McDonald, Ramdas Nagarajan, Nitya Ranganathan, Karu Sankaralingam, Simha Sethumadhavan, Premkishore Shivakumar
- TRIPS Software Team
  - Kathryn McKinley, Jim Burrill, Xia Chen, Sundeep Kushwaha, Bert Maher, Nick Nethercote, Suriya Narayanan, Sadia Sharif, Aaron Smith, Bill Yoder
- IBM Microelectronics Austin ASIC Group
- TRIPS Sponsors
  - DARPA Polymorphous Computing Architectures
  - Air Force Research Laboratories
  - National Science Foundation
  - IBM, Intel, Sun Microsystems